# WHITE PAPER

# Have you been hacked?

# Incident Handling Guidelines

# dated 5 Oct 01

**WHITE PAPER**

# Have you been hacked?
# Incident Handling Guidelines

## EXECUTIVE SUMMARY

Field Security Operations has created this document to outline the steps necessary for incident detection, response, and recovery. This paper was a follow on from the Security Workgroup that was created at the Indianapolis CDA Customer Conference. The goal of this paper is to define the “How to know when you have been compromised”, “What to do when you have”, and “What do I need to do to recover?”. The ongoing hacker threat continues to grow. It is critical that we are prepared to react in a timely fashion, to limit lost time and to ensure everything necessary is done to capture information required to catch the culprits.

**INCIDENT DETECTION**

Determining if you've been the victim of a hacking attempt can be broken down into three areas: the obvious or public-view compromise, the stealth-attack compromise, and the cover-up compromise.

The obvious or public-view compromise is just what it seems. It is the overt defacing or replacement of web pages or public download directories intended to be noticed by the public or a larger audience than just an administrator. Usually less costly and less damaging than the other types of compromises, a public-view compromise causes a lack of trust and uncertainty to permeate the user community and the customer base. This type of compromise is almost always carried out by an attack against a vulnerable server or web application. It is done to draw attention to the lack of security controls and poor system administration performed by the site owners. The administrators are usually the last to know about such compromises because the people who frequent the website are the ones that see it and inform the administrators. System Administrators generally view the system from the inside out and page defacements are harder to see this way.

Detection of this type of compromise rarely exceeds more than a few hours due to the high visibility of web pages. The more difficult process is determining which vulnerability was exploited to carry out the attack. The following areas should be checked for Windows NT Systems:

- **All temp directories.** Many times residual data is left in temp directories as strangely named files, such as ~tmp0001 or ~000001. View these files in a simple text editor, like Notepad, to see if a trail of activity is noticeable.
- **NT System Event Logs.** Examine the event logs for security violations or object-access failures or successes.
- **Application Logs.** More often than not applications keep their own logs accounting for recent activity and could possibly show footprints of the attack.
- **Security Bulletins.** Review all security bulletins relevant to your operating system and applications. Carefully trace each step of the fix or patch to ensure it was completed properly. Often times a step may be over-looked or completed improperly, and this may leave the vulnerability still exposed.
- **Check with the Vendor.** Many times vendors will be aware of potential problems and not make them public. If you notify them of a specific anomaly, they may be able to

point you in a different direction to help ascertain how the attack was carried out.

In addition to the above Windows NT examinations, the following UNIX areas should be checked:

- Examine log files for connections from unusual locations or other unusual activity. For example, look at your 'last' log, process accounting, all logs created by syslog, and other security logs. If your firewall or router writes logs to a different location than the compromised system, remember to check these logs also. ***Note:*** This is not foolproof unless you log to append-only media; many intruders edit log files in an attempt to hide their activity.
- Look for setuid and setgid files (especially setuid root files) everywhere on your system. Intruders often leave setuid copies of /bin/sh or /bin/time around to allow them root access at a later time. The UNIX find(1) program can be used to hunt for setuid and/or setgid files. For example, you can use the following commands to find setuid root files and setgid kmem files on the entire file system:

  **find / -user root -perm -4000 –print**
  **find / -group kmem -perm -2000 –print**

  ***Note: The above examples search the entire directory tree, including NFS/AFS mounted file systems. Some find(1) commands support an "-xdev" option to avoid searching those hierarchies. For example:***

  **find / -user root -perm -4000 -print -xdev**

- Another way to search for setuid files is to use the ncheck(8) command on each disk partition. For example, use the following command to search for setuid files and special devices on the disk partition /dev/rsd0g:

  **ncheck -s /dev/rsd0g**

Check your system binaries to make sure that they haven't been altered. We've seen intruders change programs on UNIX systems such as login, su, telnet, netstat, ifconfig, ls, find, du, df, libc, sync, any binaries referenced in /etc/inetd.conf, and other critical network and system programs and shared object libraries. Compare the versions on your systems with known good copies, such as those from your initial installation media. Be careful of trusting backups; your backups could also contain Trojan horses. The best way to ensure good binaries is to create a CD with binaries and add it to your incident response tool kit. This kit should be reviewed and updated with each published vulnerability.

Trojan horse programs may produce the same standard checksum and timestamp as the legitimate version. As a result, the standard UNIX sum(1) command and the timestamps associated with the programs are not sufficient to determine whether the programs have been replaced. The use of cmp(1), MD5, Tripwire, and other cryptographic checksum tools is sufficient to detect these Trojan horse programs, provided the checksum tools themselves are kept secure and are not available for modification by the intruder. Additionally, you may want to consider using a tool (i.e., PGP) to "sign" the output generated by MD5 or Tripwire, for future reference.

The stealth-attack compromise is usually carried out for purposes of gaining access to or defeating another system, not necessarily the one being compromised. One example of this would be a distributed denial-of-service attack. In this scenario, the compromised computers become zombies controlled by one or more main computers, for the purpose of attacking or compromising a third computer somewhere else. Although the compromise of the individual computers is successful, it rarely results in overt or publicly noticeable traces and often times does not compromise anything on that computer. The reason for not compromising the data or stability of the host computer is to ensure it doesn’t expose itself prior to completing its mission. This makes a stealth-attack compromise difficult to catch, at best. One thing this type of attack does do is look for allowed outbound services. It tries to communicate with the controlling host via different ports and in doing so, it determines which services and ports are allowed to originate internally and traverse the firewall. These are usually ports 80 (HTTP) and 53 (DNS). Check the following Windows NT areas:

- **Active Processes.** Review what processes are actively running on the computer and compare that list to all the known processes that are running. Often times a zombie will change the name of the rogue program to something familiar, like Explorer, and run resident under that name. Looking at this process will not be alarming unless you know that no instances of Explorer should be running. This should be done often and different times of the day.
- **System Services.** Review the stopped and running services on the system. Many of these services run hidden in the background and also run with system or administrator level privileges.
- **Scheduler Service (AT).** Check to make sure that there aren't any rogue applications scheduled to run on your system.
- **Hidden Files in NTFS Alternate Data Streams.** A stream is data that is associated with a main file or directory. Each file and directory in NTFS can have multiple data streams that are generally hidden from the user. Streams are often hidden due to the lack of stream capable

applications. Data, and even binaries, can be hidden in these streams.

- **Registry Entries.** Examine the RunOnce and Run registry keys to see if unauthorized program names are listed. This is one method of getting the program to run.
- **Autoexec.bat and Config.sys Files.** Often overlooked, these program files are still called on by Windows NT to run legacy applications. It is possible to implant code and have them run from here. Check them to ensure they are empty or do not contain unauthorized programs.
- **Review IDS and Firewall Logs.** These logs are invaluable in determining this type of compromise. Check outbound connections to non-congruent sites. In other words, if port 80 is being used for non-web traffic to a site that is not a web server, this could be an indication of an attack. Check the time of day that the packets are being sent and compare them to normal work schedules. If it is a proxy server, check the outbound log to ensure only valid protocol commands are being submitted.
- **Virus Scanning Logs.** Many zombie-like programs are detectable by virus scanners during their initial loading. Make careful review of these logs to ensure nothing was missed.
- **CPU Performance.** Monitor and check the load on the CPU, to determine if there are programs running in the background. Try to unload all known programs and monitor the activity of the CPU to see any rogue programs.

On your UNIX Systems, check for unauthorized services by doing the following:

- Inspect /etc/inetd.conf for unauthorized additions or changes. In particular, search for entries that execute a shell program (for example, /bin/sh or /bin/csh).
- Check all programs that are specified in /etc/inetd.conf, to verify that they are correct and haven't been replaced by Trojan horse programs.
- Check for legitimate services that you have commented out in your /etc/inetd.conf. Intruders may turn on a service that you previously thought you had turned off or replace the inetd program with a Trojan horse program.

The third area of compromise is the cover-up compromise. This type of attack is designed to break into a computer with the intent of retrieving data surreptitiously and covering up any sign of the retrieval. This is common in e-commerce sites when someone steals credit card information and wants it to remain unknown for as long as possible, so that the cards can be utilized to their maximum potential. This is also useful when an attacker wishes to use the trust relationship one computer has with another. Keeping the fact of the compromise hidden as long as possible is most advantageous. One down side to this attack

is the amount of time an intruder must spend on the computer. Like any thief, time is critical and the more time spent "doing the deed", the greater the risk of exposure. Often the attacker must change audit logs, logins, network statistics, permission types, access rights, and a myriad of other parameters. An administrator must have a good baseline and persistent security practices to catch this type of compromise. Check these Windows NT areas:

- **Binaries.** Check the checksum of binary files to determine if any have been corrupted or changed.
- **Examine Trust Relationships.** Look at the logs of other computers you have control of and look for anomalous behavior from the other computers – abnormal login times, what level of access obtained at each login, etc.
- **System Services.** Review the stopped and running services on the system. Many of these services run hidden in the background and also run with system or administrator level privileges.
- **Scheduler Service (AT).** Check to make sure that there aren't any rogue applications scheduled to run on your system.
- **Hidden Files in NTFS Alternate Data Streams.** A stream is data that is associated with a main file or directory. Each file and directory in NTFS can have multiple data streams that are generally hidden from the user. Streams are often hidden due to the lack of stream capable applications. Data, and even binaries, can be hidden in these streams.
- **Registry Entries.** Examine the RunOnce and Run registry keys to see if unauthorized program names are listed. This is one method of getting the program to run.
- **NTFS File Permissions.** Check for modified file permissions, especially in the %systemroot\, %systemroot\system, and %systemroot\system32 directories.
- **Examine Root-level Access.** Check for rogue accounts, locked out accounts, excessive password change requests, or changed group permissions. Any and all of these symptoms can point to a compromise. Because the attacker will seek root level access, he will be able to change logs but an account will most likely still be needed for future logins. If a normal user account is compromised, the attacker will try to exceed his authority level. If this is attempted, logs will show the attempts.

Examine all systems on the local network when searching for signs of intrusion. Most of the time, if one host has been compromised, others on the network have been compromised as well. This is especially true for networks where NIS is running or where hosts trust each other through the use of .rhosts files and/or /etc/hosts.equiv files. Also, check hosts for which your users share .rhosts access.

Check all of your systems for unauthorized use of a network-monitoring program, commonly called a sniffer or packet sniffer. Intruders may use a sniffer to capture user account and password information. For related information, see CERT advisory CA-94:01. For example, on UNIX Systems, examine all the files that are run by 'cron' and 'at'. Many intruders leave back doors in files run from 'cron' or submitted to 'at'. These techniques can allow an intruder back on the system (even after you believe you had addressed the original compromise). Also, verify that all files/programs referenced (directly or indirectly) by the 'cron' and 'at' jobs, and the job files themselves, are not world-writable.

Look everywhere on the system for unusual or hidden files (files that start with a period and are normally not shown by 'ls'), as these can be used to hide tools and information (password cracking programs, password files from other systems, etc.). A common technique on UNIX systems is to put a hidden directory in a user's account with an unusual name, something like '...' or '.. ' (dot dot space) or '..^G' (dot dot control-G). Again, the find(1) program can be used to look for hidden files, for example:

- **find / -name ".. " -print -xdev**
- **find / -name ".*" -print -xdev | cat -v**

Also, files with names such as '.xx' and '.mail' have been used (that is, files that might appear to be normal).

Examine the /etc/passwd file on the UNIX systems and check for modifications to that file. In particular, look for the unauthorized creation of new accounts, accounts with no passwords, or UID changes (especially UID 0) to existing accounts.

Check your system and network configuration files for unauthorized entries. In particular, look for '+' (plus sign) entries and inappropriate non-local host names in /etc/hosts.equiv, /etc/hosts.lpd, and in all .rhosts files (especially root, uucp, ftp, and other system accounts) on the system. These files should not be world-writable. Furthermore, confirm that these files existed prior to any intrusion and were not created by the intruder.

Remember that most hacks go unnoticed! They are in and out with logs cleaned up, all in just a few minutes. Now they have a legitimate user account with root privileges. They can do what they like and you will be none the wiser. Most people think that when they get hacked, the system crashes. The good hackers are the ones that come and go unseen. That is the danger; since everything is working, no one is looking for anything!

***Note:  All action taken during the course of an investigation should be in accordance with your organization's policies and procedures.***

## INCIDENT RESPONSE

A confirmed computer security incident should be responded to with an Incident Response and Recovery process that includes five individual procedural phases. Management and user familiarization of each phase will facilitate efficiency in the overall incident response and recovery process. These phases include:

- Preparation
- Identification
- Containment
- Eradication
- Recovery/Restoration

Documenting all steps and actions of all phases during the Incident Response and Recovery process is imperative and cannot be over emphasized.

### Preparation

Preparation is a critical phase in responding to an incident. Without proper incident reporting preparation, disorganization and confusion may result in a loss of essential data. This phase of incident handling helps to ensure that response actions are known by all and incident coordination is handled properly.

Key to preparation is the creation of written Incident Response and Recovery procedures or an Intrusion Policy Document (IPD). These procedures should be exercised periodically and would:

- Ensure the proper DOD-approved warning banner is applied to all systems. This banner clearly allows keystroke monitoring, if necessary. Keystroke information may provide a valuable record of activity and needed evidence on a compromised system.
- Provide written guidelines on how to detect and verify that a system has been compromised. Reasonable verification must be accomplished before the reporting process is to begin.
- Develop incident reporting communications requirements. This should identify the appropriate Computer Emergency Response Team (CERT) and law enforcement agency (LEA) to contact. This would include the creation of contact lists with work, home, and fax phone numbers of primary and secondary key personnel to be contacted during an incident. When all personnel are familiar with whom to call, incident reporting delay and confusion is minimized. A secure communication means should also be considered and required if warranted.

- Ensure proper backup of all essential systems and system restoration procedures. Regular backup procedures not only help to ensure operational continuity but allows administrators to check the integrity of systems and data.
- Designate an Incident Management Team. This team is key in the reporting and recovery handling of an incident and will be defined in the Identification Phase section of this paper.
- Include incident reporting and handling training to administrators and end users. This training is also essential for the designated Incident Management Team.

**Incident Identification**

The Identification Phase typically begins after a system or network abnormality has been observed and includes the confirmation and verification of the observed event. Initial indications could be received from an Intrusion Detection System (IDS) or individual system alarm. A compromise could include a number of indicators such as suspicious system or network accounting (i.e., excessive or unexplained account logon attempts, new user or system accounts, files, or file modifications or deletions). The observance of one or more of these events should stimulate a closer investigation of a possible incident. As soon as indicators confirm that a computer security incident has occurred, it should be reported immediately to the appropriate ISSO (Information Systems Security Officer).

The assignment of an Incident Management Team to be responsible for the verification, reporting and handling of the incident should be initiated at this point. The Incident Management Team should be comprised of an Incident Coordinator, and 1-4 additional team members. Incident Management Team members should be familiar with the site's organization, its mission and network architecture, as well as an in-depth knowledge of the possible compromised platform. A typical team might consist of one or all of the following:

- User - Subject matter expert for operational impact
- System Administrator - Subject matter expert for system issues
- Systems Security Analyst - Subject matter expert for system security
- Auditor - Determines economic impact of the incident.

The Incident Coordinator serves as the incident point of contact for the site and is responsible for the overall reporting and coordination process. The Incident Coordinator has the responsibility of Evidence Custodian, ensuring the documentation of the Incident Management Team's efforts while maintaining an evidence “chain of custody”. Ideally, this person would be an

individual working at the ISSM (Information Systems Security Manager) or ISSO level of information systems security management.

The Incident Management Team, led by the Incident Coordinator, will investigate and determine whether an event is an actual incident and if a compromise is confirmed, report the incident to the site's responsible regional CERT or cognizant authority. At this point the Incident Management Team will ensure that all audit trails are enabled, if not already done so. This team will take care in documenting the incident by ensuring all vital aspects of the incident are recorded, such as:

- The nature of suspicious event or incident.
- All system information (i.e., name of system and administrator, system IP, workload, etc.).
- All details relevant to the incident (i.e., the time of the incident, who and what was witnessed, etc.).
- With whom the incident or possible incident was discussed.

All reporting and coordination with outside organizations will be controlled and conducted through the Incident Coordinator. The Incident Coordinator will maintain control and handling of the incident throughout the Containment Phase of the Incident Response process.

**Incident Containment**

Incident containment should occur immediately after incident identification has occurred, or is determined to be ongoing. This phase of incident response involves limiting the scope and magnitude of the incident. The Incident Management Team will determine if the compromise is limited to one system, or if it has propagated the network, in order to determine the extent of the compromise or damage. At this point, the Incident Management Team will work closely with higher authority in evaluating the incident and determining what to do with critical information and/or computing services. It may be determined to allow a compromised system to continue to run as normal if there is a reasonable chance that the perpetrator can be exploited and the risk of damage, disruption, or compromise of data is acceptable. Or, the determination might be to remove the system from further harm by disconnecting it from the network.

It is essential that the compromised system or systems ***not*** be powered down, but rather, unplugged from the network. A power reset of a possibly compromised system may delete all possible event evidence. The team will re-confirm a proper system backup, collect, analyze, log and mark all possible evidence. It is critical that the Incident Management Team document every step of the process and that as much information about the incident be saved and documented as possible, for future forensic purposes.

The team must remember that the initial collection may escalate to supportive information or evidence of more serious violations.

**Incident Eradication**

In order to eradicate an incident, The Incident Management Team must perform an analysis on the event to determine the vulnerability that enabled the incident and any additional back doors that may have been created. Once the vulnerabilities are identified, they must be neutralized. This analysis will also reveal information used to determine how to improve defenses. Generally, host-based attacks, such as a virus incident, may simply require removing the virus through the use of virus detection/removal software or it may require a clean reformat of any disks containing infected files. Conversely, network-based incidents can be more difficult to eradicate, since any system on a network can be used to launch an attack on other addressable systems, possibly distributing the incident source.

The Incident Management Team should continue to document and coordinate all procedural steps performed in the eradication phase. Additionally, clean backups also remain very important throughout this phase of the Incident Response/Recovery process.

## INCIDENT RECOVERY/RESTORATION

The Recovery/Restoration Phase involves returning the system to full operational status. Regular system backups become very important during the Recovery/Restoration Phase. This phase consists of rebuilding the system and placing it back to operational status, while ensuring that every aspect of the system is the same as before the security incident occurred. If backups are used to rebuild the compromised system, they must be scrubbed for the eradicated event. Checks of several previous backups should be performed before restoring a backup, to assure integrity and to ensure the backup does not contain compromised or infected files. Once a system restore has been performed, it is also essential to verify that the operation was successful and that the system is back to its normal condition. Once the system has been rebuilt, validated, and returned to full production status, it must be monitored for a period of time to ensure the incident has been fully eradicated.

**REFERENCES**

For further information about the types of attack that have recently been reported to the CERT Coordination Center and for a list of new or updated files that are available for anonymous FTP, see past CERT Summaries, available in the directory:

**http://www.cert.org/summaries/**

If you suspect that your system has been compromised, please review the suggested steps in "Steps for Recovering from a UNIX Root Compromise," available from:

**http://www.cert.org/tech_tips/root_compromise.html**

Also review other appropriate files in the tech_tips directory.

To report a computer security incident to the CERT Coordination Center, please complete and return a copy of the Incident Reporting Form, available from:

**ftp://www.cert.mil/pub/forms/incident_report_form.htm**

The information on the form helps the CERT provide the best assistance, as it enables them to understand the scope of the incident, to determine if your incident may be related to any other incidents that have been reported, and to identify trends in intruder activities.

**Sources:**

CERT® Coordination Center / Intruder Detection Checklist
Logicon IASSURE Team